arXiv:1402.1047v1 [cs.DM] 5 Feb 2014

On robustly asymmetric graphs

Uriel Feige *

February 6, 2014

Abstract

O'Donnell, Wright, Wu and Zhou [SODA 2014] introduced the notion of robustly asymmetric graphs. Roughly speaking, these are graphs in which for every $0 \leq \rho \leq 1$, every permutation that permutes a ρ fraction of the vertices maps a $\Theta(\rho)$ fraction of the edges to non-edges. We show that there are graphs for which the constant hidden in the Θ notation is roughly 1.

1 Introduction

Throughout, n denotes the number of vertices in a graph and m denotes the number of edges.

Given two (labeled) graphs G and H on n vertices, their distance $dist(G, H)$ is half the Hamming distance between their adjacency matrices. If G and H have the same number of edges, then equivalently, $dist(G, H)$ is the number of pairs of vertices that are edges in G but not edges in H (or equivalently, edges in H but not edges in G). Given a permutation π and a labeled graph G, the graph resulting from applying π to the vertices of G is denoted by G_π. Given a set of n elements, a k-permutation is a permutation that permutes k of the elements and leaves exactly $n - k$ elements in place.

Definition 1 *For $\delta > 0$, a graph $G(V, E)$ with n vertices and m edges is δ-asymmetric if for every $1 \leq k \leq n$ and every k-permutation π,*

$$dist(G, G_\pi) \geq \delta \frac{k}{n} m.$$

*Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot 76100, Israel. E-mail: uriel.feige@weizmann.ac.il. Work supported in part by The Israel Science Foundation (grant No. 621/12) and by the I-CORE Program of the Planning and Budgeting Committee and The Israel Science Foundation (grant No. 4/11).

In our theorems concerning asymmetric graphs we shall only consider graphs with at most $\frac{1}{2}\binom{n}{2}$ edges. By taking their complements one can derive obvious consequences regarding graphs with more than $\frac{1}{2}\binom{n}{2}$ edges, but these consequences are omitted from this manuscript.

Our main theorem is the following.

Theorem 2 *For every $\delta < 1$ there is some constant $c > 0$, such that for every n and for every $cn \leq m \leq n^2/c$, there is a graph G with n vertices and m edges that is δ-asymmetric.*

Theorem 2 is best possible in the following sense. Taking an arbitrary k-permutation with $k = n$ implies that graphs cannot be δ-asymmetric for $\delta > 1$, and not even for $\delta = 1$ (as there are n-permutations that preserve some edges). A value of $\delta = 1 - o(1)$ (where the $o(1)$ term tends to 0 when n and c grow) is the best that one can hope for, and indeed, our proofs indeed provide such bounds. (One can formulate Theorem 2 with both δ and c being functions of n, though for simplicity we have not done so.) However, for small values of k, the value of δ that can conceivably be achieved is not just 1, but in fact arbitrarily close to 2 (because an ϵ fraction of the vertices may be incident with nearly a 2ϵ fraction of the edges). Indeed, *our proofs provide this improved bound*. The requirement that $m \geq cn$ is unavoidable because every regular graph has an n-permutation that preserves $\Omega(n)$ edges. The requirement that $m \leq n^2/c$ is unavoidable because a random permutation will preserve roughly $m\frac{m}{\binom{n}{2}} \simeq 2m/c$ edges in expectation.

Our proof of Theorem 2 is based on the probabilistic method. That is, we do not explicitly construct graphs satisfying Theorem 2, but merely show their existence. In our proof, we consider two different ranges of values for m, and for each of them show a different randomized construction. These two ranges are handled by Theorem 3 and by Theorem 4, and the combination of these two theorems implies Theorem 2. We now provide more details of these two randomized constructions.

Recall the random graph model $G_{n,p}$ of Erdos and Renyi in which every edge is included independently with probability p.

Theorem 3 *For a sufficiently large constant c, for every p satisfying $\frac{c \log n}{n} \leq p \leq \frac{1}{2}$, the random graph $G_{n,p}$ is δ-asymmetric almost surely. Starting at the low end of probabilities for p, the value of δ tends to 1 as c grows, and then at the high end (when p is constant), $\delta \simeq 1 - p$.*

Theorem 3 does not prove the existence of δ-asymmetric graphs with a linear number of edges. The problem is that when $p = O(1/n)$, a graph

sampled in the $G_{n,p}$ model is likely to have isolated vertices, and these can be permuted without changing the graph. Likewise, the graph is likely to have pairs of degree one vertices connected to the same vertex (or to each other), and the two vertices within such a pair can be permuted as well. The 2-core of $G_{n,p}$ may be asymmetric, but is unlikely to be δ-asymmetric for constant values of δ. It is likely to have chains of degree 2 vertices of super-constant length, and flipping such a chain permutes a super-constant number of vertices, while increasing the Hamming distance only by a constant.

To exhibit δ-asymmetric graphs with a linear number of edges, one can consider random regular graphs. However, for simplicity of the proofs, we consider a model that we call $G_{n,p,d}$ that extends the $G_{n,p}$ model in the following way. Given a graph sampled from $G_{n,p}$, iteratively consider the vertices of G, where vertex i is considered in iteration i. Let d_i be the degree of vertex i at that point. If $d_i \geq d$, do nothing at that iteration. If $d_i < d$, add $d - d_i$ *auxiliary edges* incident with i, connecting it to vertices who were not previously neighbors of i, chosen uniformly at random.

Theorem 4 *For a sufficiently large constant c, let $\frac{c}{n} \leq p \leq \frac{(\log n)^2}{n}$, and let $d = \lceil p(n-1) \rceil$. Then the random graph sampled from $G_{n,p,d}$ is δ-asymmetric almost surely, for $\delta = 1 - O(1/\sqrt{d})$.*

2 Related work

A graph is said to be asymmetric if it has no nontrivial automorphism. Erdos and Renyi [3] proved that a graph sampled at random from the $G_{n,p}$ model is asymmetric with high probability, for p in the range $\frac{\ln n}{n} < p < 1 - \frac{\ln n}{n}$. Later work [2, 5, 4] established that random d-regular n-vertex graphs are asymmetric with high probability, for d in the range $3 \leq d \leq n - 4$.

The notion of robustly asymmetric graphs (under a slightly different definition, see the last sentence of this paragraph) was introduced by O'Donnell, Wright, Wu and Zhou, as part of work that studied the computational complexity of the *robust graph isomorphism* problem [6]. They prove a theorem similar to our Theorem 3, but with a value of $\delta = \frac{1}{240}$. They comment that they did not work hard to optimize the constants in the theorem statement, and that it is interesting to explore the limits of these constants. Our Theorem 3 achieves the best possible constant of $\delta \simeq 1$. (See also the discussion following Theorem 2.) For graphs of constant average degree, the authors of [6] suggest the study random d-regular graphs as a future research direction. Our Theorem 4 considers a somewhat different distribution over

random graphs that is presumably easier to analyze, and gives the type of results that one would hope to prove for random d-regular graphs. The paper [6] also contains results for $G_{n,p}$ with $p = O(1/n)$ (hence constant average degree), but these results refer to a version of robust asymmetry which disregards automorphisms that permute only a small fraction of the vertices.

2.1 Notation and preliminaries

For disjoint sets S, T of vertices in a graph G, $E(S)$ denotes the set of edges induced by S, and $E(S,T)$ denotes the set of edges with one endpoint in S and the other in T. The notation $E[x]$ denotes the expectation of a random variable x.

The reader is assumed to be familiar with standard probabilistic reasoning, and with large deviation bounds such as those of Chernoff and Azuma. See [1] for example.

3 Proofs

We first prove Theorem 3.

Proof. Fix $\epsilon > 0$ to be a small constant. All results will hold when $n > n_0$, where n_0 may depend on ϵ. In particular, we shall take $n_0 > 1/\epsilon$.

Consider a random graph $G(V,E)$ selected at random from $G_{n,p}$. Consider an arbitrary set $S \subset V$ of vertices and let $k = |S|$. We say that an edge $e \in E$ is *covered* by S if at least one of its endpoints is in S. Let m_S denote the number of edges covered by S. Hence $m_S = |E(S, V \setminus S)| + |E(S)|$.

Proposition 5 *With probability at least* $1 - \frac{\epsilon}{n\binom{n}{k}}$ *over the choice of* G,

$$|m_S - p\left(\binom{k}{2} + k(n-k)\right)| \leq O\left(\sqrt{pkn \log \binom{n+1}{k}}\right)$$

Proof. Fixing S, the expectation of m_S under the a random choice of a graph from $G_{n,p}$ is $E[m_S] = p\left(\binom{k}{2} + k(n-k)\right)$. Observe that this value is $\Theta(pkn)$, and for our choice of p in the statement of Theorem 3, it is at least $\log n$. As the choices of edges in the $G_{n,p}$ model are independent, we can apply standard bounds on large deviations for sums of independent random variables (such as the Chernoff bound) to deduce that the probability of a deviation by a factor of $a\sqrt{pkn}$ from the expectation is exponentially small

in a^2. Plugging $a = \sqrt{\log(n\binom{n}{k}/\epsilon)}$, and observing that for our range of parameters $\log(\binom{n+1}{k}) = \Omega\left(\log(n\binom{n}{k}/\epsilon)\right)$, the proposition follows. (We used that assumption that $n \geq 1/\epsilon$. The change from $\binom{n}{k}$ to $\binom{n+1}{k}$ was made so as to address the case that $k = n$.) ■

We say that S is *typical* if m_S is in the range specified by Proposition 5. Observe that by a union bound over all choices of S, it is likely that all $S \subset V$ are typical.

Lemma 6 *Given S of size $|S| = k$, let π be an arbitrary k-permutation that moves every vertex within S but keeps every vertex in $V \setminus S$ fixed. Then conditioned on S being typical, with probability at least $1 - \frac{\epsilon}{n^{k+1}}$ over the choice of G,*

$$dist(G, G_\pi) \geq m_S \frac{\binom{k}{2} + k(n-k) - m_S}{\binom{k}{2} + k(n-k)} - O\left(\sqrt{m_S k \log n}\right)$$

Proof. Let P denote the set of unordered pairs of vertices such that either both vertices are from S, or one from S and one from $V \setminus S$. Hence $|P| = \binom{k}{2} + k(n-k)$. Each member of P may potentially be an edge in G. Consider now how the permutation π acts on P. It may have a fixpoint with respect to P: if π maps u to v and maps v to u, then the unordered pair (u, v) is a fixpoint. Let f denote the number of fixed points, and note that necessarily $f \leq k/2 \leq |P|/(n+1)$.

Conditioned on the value of m_S, for a random choice of G, the actual edges covered by S are distributed uniformly at random over the members of P. Let x_i be an indicator random variable for the event that π maps the vertex pair of the ith edge covered by S to a vertex pair that is not an edge in P. Denote $X = \sum_{i=1}^{m_S} x_i$, and observe that $dist(G, G_\pi) = X$. Hence it remains to estimate X.

We claim that $Pr[x_i = 1] = \frac{|P|-f}{|P|} \frac{|P|-m_S}{|P|-1}$. The first term corresponds to the ith edge not being a fixpoint of π, whereas the second term corresponds to none of the other $m_S - 1$ edges covered by S being the image of edge i under π. Hence the expectation of X satisfies

$$E[X] = E[\sum_{i=1}^{m_S} x_i] = m_s \frac{|P| - f}{|P|} \frac{|P| - m_S}{|P| - 1} \geq \frac{n-2}{n-1} \frac{m_S(|P| - m_S)}{|P|}$$

A martingale argument proves that X is concentrated around its mean. Fixing π and exposing the m_f covered edges one by one, every new edge

can cause X to increase by at most 1 (by the edge itself), and to decrease by at most 1 (by being the image under π of a previous edge). Hence the martingale sequence enjoys the bounded difference property. As $E[X] = \Theta(m_S)$, the probability of deviating by a factor of $a\sqrt{m_S}$ is exponentially small in a^2 (by Azuma's inequality). Plugging $a = \sqrt{\log(n\binom{n}{k}/\epsilon)}$ the lemma follows. (The term n/ϵ that should be present inside the log disappears due to the O notation and the assumption that $n \geq 1/\epsilon$. The term $\frac{n-2}{n-1}$ can be omitted because its contribution is negligible compared to the standard deviation.) ■

We now complete the proof of Theorem 3. By the union bound, every set S is typical, and then another union bound shows that for all permutations the distances are as in Lemma 6. For $p \geq \frac{c \log n}{n}$ the error term in Proposition 5 is small compared to the main term, and we have that for every S it holds that $m_S \simeq p|P|$ (where P is as in Lemma 6). Likewise, in Lemma 6 the error term can also be seen to be small compared to the main term, and then $dist(G, G_\pi) \simeq p(1-p)|P|$. The ratio between $|P|$ and k is $n - \frac{k+1}{2}$ and is minimized when $k = n$, giving roughly $n/2$. As $m \simeq \frac{pn^2}{2}$, we get that:

$$dist(G, G_\pi) \simeq p(1-p)|P| \geq p(1-p)\frac{kn}{2} \simeq (1-p)\frac{k}{n}m$$

establishing that $\delta \simeq 1 - p$. ■

Before proving Theorem 4, we establish some properties of the $G_{n,p,d}$ model.

In the $G_{n,p}$ model, every edge is present independently with probability p. The following proposition provides an extension of this property to the $G_{n,p,d}$ model.

Proposition 7 *Let F be an arbitrary (possibly empty) set of edges, and let $e \notin F$ be an arbitrary edge. Then in the $G_{n,p,d}$ model, the probability that e is present conditioned on F being present satisfies*

$$p \leq Pr[e \in E | F \subset E] \leq p + \frac{2d}{n-1}$$

Proof. Let u and v denote the endpoints of e (namely, $e = (u, v)$). The probability that e is present not as an auxiliary edge (namely, as one of the original edges of $G_{n,p}$) is exactly p, regardless of F. The information conveyed through the event $F \subset E$ may only affect the probability the e

is chosen as an auxiliary edge. However, because we condition on $F \subset E$ (rather than say on $F \cap E$ being empty), the probability that u chooses e as an auxiliary edge is at most $\frac{d}{n-1}$, and the same applies to v. ■

Lemma 8 *Let G be a graph selected at random from $G_{n,p,d}$, where $\frac{c}{n} \leq p \leq \frac{log^2 n}{n}$ and $d = \lceil (p-1)n \rceil$, where c is a sufficiently large constant and n is sufficiently large. Then G satisfies the following properties almost surely:*

1. *The average degree in G is $pn + O(\sqrt{pn}) = d + O(\sqrt{d})$.*
2. *Every two vertices of G have at most two common neighbors.*
3. *Every set S of size at most n/d^2 induces at most $3|S|$ edges.*

Proof. We prove the three properties in the order stated in the lemma.

1. The expected degree of a vertex in $G_{n,p}$ is $p(n-1)$ and its variance is of the same order. Hence in expectation, the number of auxiliary edges added per vertex is $O(\sqrt{pn})$, and the total expected number of auxiliary edges in $O(n\sqrt{pn})$. The actual number is concentrated around its expectation (e.g., by a Martingale argument). Hence the auxiliary edges contribute $O(\sqrt{pn})$ to the average degree.

2. There are $\binom{n}{2}\binom{n-2}{3}$ possible ways of choosing two vertices and three potential common neighbors. The probability that all required six edges are present is roughly at most $(3p)^6$, by Proposition 7. As $p^6 n^5 = o(1)$ for our choice of p, a union bound shows that almost surely no two vertices share three common neighbors

3. Using Proposition 7, the probability that S induces M edges is at most $\binom{|S|^2/2}{M}(3p)^M$. As there are $\binom{n}{\ell}$ choices of S of cardinality ℓ, the property fails to hold with probability at most

$$\begin{aligned}
\sum_{\ell=1}^{n/d^2} \binom{n}{\ell}\binom{\ell^2/2}{3\ell}(3p)^{3\ell} &\simeq \sum_{\ell=1}^{n/d^2} \left(\frac{en}{\ell}\right)^{\ell} \left(\frac{e\ell}{6}\right)^{3\ell} \left(\frac{3d}{n}\right)^{3\ell} \\
&\leq \sum_{\ell=1}^{\sqrt{n}} \frac{e^{4\ell} d^{\ell}}{2^{3\ell} n^{\ell}} + \sum_{\ell=\sqrt{n}}^{n/d^2} \frac{e^{4\ell}}{2^{3\ell} d^{\ell}} \\
&= o(1).
\end{aligned}$$

■

We can now prove Theorem 4.

Proof. We handle separately two ranges of values for k.

Large k. This covers the range $n/d^2 < k \leq n$. One can verify that the proof of Theorem 3 with some straightforward modifications applies in this case.

Small k. This covers the range $1 < k \leq n/d^2$. Let S be the set of size k permuted by π. Edges in $E[S, V \setminus S]$ are referred to as *outgoing*. By construction, every vertex of S has degree at least d. By item 3 of Lemma 8, S has at least $(d-6)|S|$ outgoing edges. By item 2 of Lemma 8, at most two outgoing edges per vertex are preserved by π. Hence π maps at least $(d-8)k$ edges to non-edges. By item 1 of Lemma 8, the total number of edges in G is at most $\frac{dn}{2} + O(n\sqrt{d})$. Hence the distance between G and G_π is $(2 - O(1/\sqrt{d}))\frac{k}{n}m$, giving a value of δ which is even larger than claimed in the theorem.

■

References

[1] Noga Alon, Joel Spencer: The Probabilistic Method. John Wiley 1992. (Later editions also available.)

[2] Bela Bollobas. The asymptotic number of unlabelled regular graphs. The Journal of the London Mathematical Society, Second Series, 26(2):201-206, 1982.

[3] Paul Erdos and Alfred Renyi. Asymmetric graphs. Acta Mathematica Hungarica, 14(3):295-315, 1963.

[4] Jeong Han Kim, Benny Sudakov, Van Vu. On the asymmetry of random regular graphs and random graphs. Random Structures and Algorithms, 21(3-4):216-224, 2002.

[5] Brendan McKay and Nicholas C. Wormald. Automorphisms of random graphs with specified vertices. Combinatorica, 4(4):325-338, 1984.

[6] Ryan O'Donnell, John Wright, Chenggang Wu, Yuan Zhou. Hardness of robust graph isomorphism, Lasserre gaps, and asymmetry of random graphs. SODA 2014. http://arxiv.org/abs/1401.2436.